

INTELLIGENCE BRIEFING

Security Command Center

TLP:CLEAR

2026-07-02 07:12 UTC

# LLM Hallucinated Domains Open a New Supply Chain Attack Lane: What Security Teams Must Do Now

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0314
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Organizations using LLMs for development assistance, code generation, and package recommendations; no specific product versions identified
Published	2026-07-01T11:17:14
Discovery Source	Rss

## Executive Summary

Large language models used in developer workflows have been observed hallucinating plausible-sounding but nonexistent package names and web domains. Researchers have identified a technique, reported as 'phantom squatting,' in which attackers pre-register those fabricated identifiers to serve malware or conduct dependency confusion attacks, bypassing conventional domain-monitoring tools because the squatted names bear no resemblance to any legitimate brand. While confirmed exploitation at scale has not been established by the reviewed source material, the attack surface is broad: any organization where developers act on LLM-generated package or dependency recommendations without independent verification is potentially exposed to a novel supply chain integrity failure.

## Technical Analysis

The phantom squatting technique exploits a structural property of LLMs: when asked to recommend libraries, packages, or resources, models occasionally generate plausible-sounding names that do not correspond to real software. Unlike typosquatting, which relies on visual similarity to known packages, phantom squatting produces names with no legitimate counterpart, meaning conventional domain-monitoring and brand-similarity tooling has no anchor to flag them against.

According to reporting in Dark Reading, attackers can pre-register these hallucinated identifiers across package registries (npm, PyPI, RubyGems, and others) and domain name spaces before any developer encounters them. When a developer queries an LLM for a dependency recommendation, receives a hallucinated package

name, and installs it without verification, the malicious artifact resolves and executes in the development or build environment.

The attack chain maps directly to several MITRE ATT&CK techniques. Infrastructure acquisition (T1583.001) and obtaining tooling (T1588.001) describe the pre-registration phase. Supply chain compromise at the dependency level (T1195.001) captures the delivery mechanism. User execution via malicious links and files (T1204.001, T1204.002) describes the developer action that triggers the chain. Spearphishing via link (T1566.002) and standard application layer protocol communications (T1071.001) describe potential post-installation command-and-control or data exfiltration behaviors.

The relevant CWEs, CWE-829 (Inclusion of Functionality from Untrusted Control Sphere), CWE-494 (Download of Code Without Integrity Check), and CWE-345 (Insufficient Verification of Data Authenticity), all point to the same defensive gap: organizations lack controls that require cryptographic or registry-authoritative verification before a dependency is resolved and executed.

The OWASP LLM Top 10 framework, referenced in secondary sources, explicitly identifies supply chain risks as a top-tier LLM security concern for 2026, and vendor-tier sources from Palo Alto Networks and Wiz both document dependency confusion and prompt-adjacent supply chain attacks as active LLM risk categories.

No CVE identifier applies because this is an emergent attack class, not a discrete software vulnerability. The source quality score of 0.64 and the absence of confirmed large-scale exploitation mean claims about active campaigns should be treated as medium-confidence. The conceptual threat mechanism is documented in OWASP LLM Top 10 and vendor research; the operational prevalence and confirmed exploitation at scale are not yet established.

## Action Checklist

1. Step 1: Assess exposure, audit developer workflows to determine whether LLMs (including GitHub Copilot, ChatGPT, Gemini, or similar tools) are used to generate package names, dependency lists, or installation commands; document which teams and pipelines are affected
2. Step 2: Review controls, verify that all package installations in CI/CD pipelines require checksum or cryptographic integrity verification before execution (CWE-494 mitigation); confirm that dependency resolution is pinned to approved registry mirrors and that CIS Safeguard 2.1 (Establish and Maintain a Software Inventory) is enforced across build environments
3. Step 3: Enforce dependency verification policy, implement or audit controls aligned with NIST SI-7 (Software, Firmware, and Information Integrity) and CM-5 (Access Restrictions for Change) to restrict outbound package resolution to approved registries only; ensure that unapproved or unrecognized package names trigger a hold-and-review step rather than silent installation
4. Step 4: Update threat model, add 'LLM-hallucinated dependency injection' as a named attack pattern in your supply chain threat register; map it to T1195.001 (Supply Chain Compromise: Compromise Software Dependencies and Development Tools) and document phantom squatting as a variant distinct from typosquatting
5. Step 5: Train developers, brief engineering teams on the phantom squatting mechanism; establish a policy requiring independent verification of any package name sourced from an LLM recommendation against the official registry before installation
6. Step 6: Monitor developments, track follow-up research, registry-level detections, and any threat intelligence feeds that begin publishing hallucinated package name indicators; consult the Dark Reading

source article for updated indicators as the threat matures

## IR / Forensic Enrichment

<b>Triage Priority</b>	URGENT
<b>Escalation Criteria</b>	Escalate immediately to CISO and legal if forensic review of CI/CD build logs or developer workstations confirms that a phantom-squatted package was successfully installed and executed in any environment — particularly if the affected pipeline had access to secrets, signing keys, production credentials, or customer data, which would trigger breach notification assessment under applicable regulations.
<b>Recovery Notes</b>	After enforcing registry allowlisting and dependency pinning, rebuild all CI/CD build agents from a known-good image and re-run all pipeline builds against the cleaned, hash-pinned lockfiles to ensure no phantom-squatted packages persist in build caches or artifact repositories. For any package confirmed to have been installed from a hallucinated or phantom-squatted source, treat the build output as potentially compromised — audit all artifacts produced by that pipeline for signs of supply chain tampering (unexpected binary sections, anomalous outbound connections at runtime, or code not present in the source repository). Maintain enhanced monitoring of outbound DNS and network connections from build agents and production systems for 90 days post-remediation, specifically watching for beaconing to newly registered domains consistent with phantom-squatted package callback infrastructure.
<b>Forensic Artifacts</b>	CI/CD pipeline execution logs (GitHub Actions logs, GitLab CI job logs, Jenkins console output) showing exact package names resolved and downloaded during builds in the 90 days prior to detection — hallucinated package names will appear as successful or failed resolution attempts against registries   DNS query logs from build agents and developer workstations filtered for queries to package registry domains (pypi.org, registry.npmjs.org, pkg.go.dev) cross-referenced against a WHOIS/registration-date lookup for each resolved package name — phantom-squatted domains will show registration dates clustered around LLM hallucination research publication dates   Package manager cache directories on build agents and developer workstations (~/.cache/pip/, ~/.npm/_cacache/, \$GOPATH/pkg/mod/) containing downloaded tarballs — these can be hashed and compared against official registry checksums to detect tampered packages delivered from phantom-squatted sources   Network flow or proxy logs showing outbound HTTP/HTTPS connections from build agents to package registry endpoints, including full URLs with package names and versions, source IP, timestamp, and bytes transferred — a phantom-squatted package fetch will show a connection to the legitimate registry domain but downloading a package with an LLM-plausible fabricated name   Developer IDE extension logs and LLM chat export history (VS Code Copilot logs at platform-specific log paths, exported ChatGPT or Gemini conversation history) containing the original LLM-generated package name suggestions — these establish the causal link between the hallucination event and any subsequent installation attempt and are critical for scoping which developers and projects were exposed

### Per-Action IR Details

**Step 1: Assess exposure — audit developer workflows to determine whether LLMs (including GitHub Copilot, ChatGPT, Gemini, or similar tools) are used to generate package names, dependency lists, or installation commands; document which teams and pipelines are affected**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: Establishing IR capability and identifying assets, workflows, and exposure surface before an incident occurs

**Controls:** CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 2.1 (Establish and Maintain a Software Inventory), NIST AC-1 (Policy And Procedures)

**Compensating:** Run 'pip list --format=freeze > requirements\_audit.txt' and 'npm list --all > npm\_audit.txt' across all build environments, then diff against your approved package baseline. For CI/CD pipelines using GitHub Actions or GitLab CI, grep workflow YAML files for 'pip install', 'npm install', or 'go get' commands that reference dynamically generated package names rather than pinned lockfiles. A two-person team can complete this with bash: 'grep -rn "pip install|npm install|go get" .github/workflows/ --include="\*.yml"'.

**Evidence:** This step does not alter live state. Capture before proceeding: export current lockfiles (package-lock.json, requirements.txt, go.sum, Pipfile.lock) from all active build agents and developer workstations to establish a baseline of what packages were installed and when. Also export CI/CD pipeline execution logs showing which package names were resolved during recent builds — these logs may contain hallucinated package names that were already silently installed.

## **Step 2: Review controls — verify that all package installations in CI/CD pipelines require checksum or cryptographic integrity verification before execution (CWE-494 mitigation); confirm that dependency resolution is pinned to approved registry mirrors and that CIS Safeguard 2.1 (Establish and Maintain a Software Inventory) is enforced across build environments**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: Establishing preventive controls and verifying their effectiveness prior to exploitation

**Controls:** CIS 2.1 (Establish and Maintain a Software Inventory), CIS 2.2 (Ensure Authorized Software is Currently Supported), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), NIST SI-2 (Flaw Remediation)

**Compensating:** For Python pipelines, enforce hash pinning in requirements.txt using 'pip-compile --generate-hashes' from pip-tools (free). For Node.js, enforce 'npm ci' instead of 'npm install' — 'npm ci' requires a committed package-lock.json and fails if the lock is absent or mismatched. For Go modules, verify 'go.sum' is committed and enable GONOSUMCHECK is not set. Configure Artifactory OSS or Nexus Repository OSS (both free tiers available) as a proxy registry mirror, then block direct outbound access to pypi.org, registry.npmjs.org, and pkg.go.dev via host-based firewall rules ('iptables -A OUTPUT -d pypi.org -j DROP') so all resolution flows through the approved mirror.

**Evidence:** This step does not alter live state. Before changing any registry configuration, export the current pip.conf, .npmrc, and GOPROXY environment variable values from all build agents — these document whether outbound registry access was previously unrestricted and will be needed to demonstrate the pre-remediation exposure window in any post-incident review.

## **Step 3: Enforce dependency verification policy — implement or audit controls aligned with NIST AC-4 (Information Flow Enforcement) to restrict outbound package resolution to approved registries only; ensure that unapproved or unrecognized package names trigger a hold-and-review step rather than silent installation**

**NIST Phase:** Containment

**Reference:** NIST 800-61r3 §3.3 — Containment Strategy: Restricting the attack surface to prevent further ingestion of malicious phantom-squatted packages while investigation proceeds

**Controls:** NIST AC-4 (Information Flow Enforcement), CIS 4.4 (Implement and Manage a Firewall on Servers), CIS 4.2 (Establish and Maintain a Secure Configuration Process for Network Infrastructure)

**Compensating:** Implement an allowlist-only egress firewall rule on all CI/CD build agents restricting outbound TCP 443 to the IP ranges of your approved registry mirror only. On Linux build agents: 'iptables -P OUTPUT DROP && iptables -A OUTPUT -d j ACCEPT && iptables -A OUTPUT -d -j ACCEPT'. For GitHub Actions, use a third-party network egress restriction action (e.g., step-security/harden-runner, open source) to enforce allowed endpoints. Add a pre-install script hook that queries the approved package manifest and exits non-zero if the package name is absent — this forces a pipeline failure and human review before any unrecognized package is fetched.

**Evidence:** Before enforcing egress restrictions, capture: (1) DNS query logs from build agents for the 30 days prior to enforcement — hallucinated package names will appear as DNS lookups for nonexistent or recently registered domains; (2) network flow logs showing outbound connections from build agents to package registries, noting any domains registered within the past 90 days which may indicate pre-registered phantom-squatted names; (3) proxy or firewall logs showing package download events with full URLs, timestamps, and resolved IP addresses.

**Step 4: Update threat model — add 'LLM-hallucinated dependency injection' as a named attack pattern in your supply chain threat register; map it to T1195.001 (Supply Chain Compromise: Compromise Software Dependencies and Development Tools) and document phantom squatting as a variant distinct from typosquatting**

**NIST Phase:** Post Incident

**Reference:** NIST 800-61r3 §4 — Post-Incident Activity: Updating policies, threat models, and organizational knowledge to reflect newly identified attack patterns

**Controls:** NIST RA-3 (Risk Assessment), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

**Compensating:** Add a MITRE ATT&CK Navigator layer (free, browser-based at [attack.mitre.org/resources/attack-navigator/](https://attack.mitre.org/resources/attack-navigator/)) annotating T1195.001 with a custom note describing the phantom squatting variant and your organization's specific LLM tooling exposure (e.g., 'GitHub Copilot used in 3 teams; ChatGPT used ad hoc by developers for package recommendations'). Store the navigator layer JSON in your threat register git repository alongside the threat model document so it is version-controlled and auditable.

**Evidence:** This step does not alter live state and does not require volatile capture. For completeness, document any LLM chat session logs or IDE Copilot suggestion history that may have been retained — some IDEs log Copilot suggestions locally (e.g., VS Code extension logs at '%APPDATA%\Code\logs\' on Windows or '~/.config/Code/logs/' on Linux), which could provide examples of hallucinated package names suggested to your developers.

**Step 5: Train developers — brief engineering teams on the phantom squatting mechanism; establish a policy requiring independent verification of any package name sourced from an LLM recommendation against the official registry before installation**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: Building organizational awareness and human-layer controls to reduce susceptibility to novel attack vectors that bypass automated tooling

**Controls:** NIST AC-1 (Policy And Procedures), CIS 6.1 (Establish an Access Granting Process)

**Compensating:** Distribute a one-page developer reference card (internal wiki page or printed) listing the exact verification steps: (1) take the package name from the LLM output, (2) search [pypi.org/project/](https://pypi.org/project/), [npmjs.com/package/](https://npmjs.com/package/), or [pkg.go.dev/](https://pkg.go.dev/) directly in a browser — do not use the LLM to confirm its own suggestion, (3) check the package's creation date on the registry — a package registered within the last 90 days with no prior release history and a name matching an LLM-plausible pattern is high-risk, (4) verify the package has a credible maintainer history and download count before installing. Post a Slack or Teams reminder in developer channels each time a new phantom squatting incident is reported publicly.

**Evidence:** This step does not alter live state. No volatile capture is required. Document the training delivery (attendance records, wiki page version history, or Slack/Teams message timestamps) to support audit evidence for NIST AC-1 policy dissemination requirements.

**Step 6: Monitor developments — track follow-up research, registry-level detections, and any threat intelligence feeds that begin publishing hallucinated package name indicators; consult the Dark Reading source article for updated indicators as the threat matures**

**NIST Phase:** Post Incident

**Reference:** NIST 800-61r3 §4 — Post-Incident Activity: Sharing intelligence, monitoring for emerging indicators, and integrating threat intelligence into ongoing detection posture

**Controls:** NIST AU-6 (Audit Record Review, Analysis, And Reporting), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

**Compensating:** Subscribe to OSV (osv.dev, free) and Socket.dev (free tier) which both monitor npm and PyPI for newly published packages exhibiting suspicious signals (new maintainer, no prior history, name resembling common LLM-generated patterns). Set up an RSS feed or GitHub watch on the Checkmarx Supply Chain Security research repository and the OSSF (Open Source Security Foundation) blog. For internal detection, write a daily cron job that queries the PyPI JSON API for any package name found in your codebase that was registered within the last 180 days: `'curl https://pypi.org/pypi/|jq .info.requires_python,.info.author,.urls[0].upload_time'` — flag packages with upload dates post-dating their first appearance in your lockfiles.

**Evidence:** This step does not alter live state. Retain any threat intelligence reports, feed exports, or registry query results as dated artifacts in your threat register. If a specific hallucinated package name is published as an indicator of compromise by a researcher or registry, immediately cross-reference it against your historical CI/CD build logs and developer workstation package lists captured during Step 1 to determine retrospective exposure.

## Detection Guidance

No discrete, verifiable IOCs (domains, hashes, package names) were present in the source material reviewed. The Dark Reading article cited in the sources may contain specific hallucinated package name examples or registry artifacts, consult that source directly for any published indicators.

Behavioral patterns and log sources to hunt against:

1. Package manager logs (npm, pip, gem, cargo): Flag any package installation that resolves successfully but whose registry entry was created recently (within 30-90 days) and has zero prior download history or no associated source repository.
2. CI/CD pipeline logs: Alert on dependency resolution requests that do not match an allowlisted or previously approved package manifest. Deviations from a pinned lockfile (package-lock.json, requirements.txt, Pipfile.lock) warrant review.
3. DNS and proxy logs: Unusual outbound connections to newly registered domains during build or install phases, particularly domains with no historical resolution record, align with T1583.001 infrastructure acquisition and should be triaged.
4. Developer workstation telemetry: File execution events following package installation from an unrecognized source map to T1204.002 (Malicious File execution). EDR telemetry showing post-install script execution (postinstall hooks in npm, setup.py in PyPI) from packages not in an approved inventory should trigger review.
5. Audit logging (NIST AU-2, AU-6): Ensure build environment audit logs capture package resolution events with sufficient detail (package name, registry endpoint, resolved version, hash) to support after-the-fact investigation.

Gap audit: Verify that your software composition analysis (SCA) tooling checks package names against known-good registry baselines, not only against known-malicious signatures. Phantom-squatted packages will not appear in threat intelligence feeds until after discovery, detection must rely on behavioral and provenance signals, not signature matching.

## Framework Mappings

### MITRE-ATTACK

- **T1583.001** — Domains
- **T1204.001** — Malicious Link
- **T1071.001** — Web Protocols

- **T1204.002** — Malicious File
- **T1588.001** — Malware
- **T1195.001** — Compromise Software Dependencies and Development Tools
- **T1566.002** — Spearphishing Link

#### NIST-800-53R5

- **AT-2** — Literacy Training and Awareness
- **SC-7** — Boundary Protection
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SI-8** — Spam Protection
- **SI-7** — Software, Firmware, and Information Integrity
- **CM-3** — Configuration Change Control
- **SR-2** — Supply Chain Risk Management Plan

#### OWASP-TOP10-2021

- **A08:2021** — Software and Data Integrity Failures

#### CIS-V8

- **2.5** — Allowlist Authorized Software
- **2.6** — Allowlist Authorized Libraries
- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks
- **15.1** — Establish and Maintain an Inventory of Service Providers
- **8.2** — Collect Audit Logs

#### ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information
- **A.5.21** — Managing information security in the ICT supply chain

#### HIPAA-SECURITY

- **164.308(a)(5)(i)** — Security Awareness and Training

#### NIST-CSF-2

- **GV.SC-01** — Cybersecurity supply chain risk management program
- **DE.CM-01** — Networks and network services are monitored

#### SOC2-TSC

- **CC9.2** — Manages risks associated with vendors and business partners

## MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1583.001	Domains	Resource-Development
T1204.001	Malicious Link	Execution
T1071.001	Web Protocols	Command-And-Control
T1204.002	Malicious File	Execution
T1588.001	Malware	Resource-Development
T1195.001	Compromise Software Dependencies and Development Tools	Initial-Access
T1566.002	Spearphishing Link	Initial-Access

## Sources

Source	URL	Tier
<b>Security News</b>	<a href="https://www.darkreading.com/endpoint-security/phantom-squatting-ai-...">https://www.darkreading.com/endpoint-security/phantom-squatting-ai-...</a>	T2
<b>OWASP LLM Top 10: AI Security Risks to Know in 2026</b>	<a href="https://elevateconsult.com/insights/owasp-llm-top-10-security-vulne...">https://elevateconsult.com/insights/owasp-llm-top-10-security-vulne...</a>	T3
<b>LLM Security in 2025: Risks, Examples, and Best Practices</b>	<a href="https://www.oligo.security/academy/llm-security-in-2025-risks-examp...">https://www.oligo.security/academy/llm-security-in-2025-risks-examp...</a>	T3
<b>What Is LLM (Large Language Model) Security?   Starter Guide</b>	<a href="https://www.paloaltonetworks.com/cyberpedia/what-is-llm-security">https://www.paloaltonetworks.com/cyberpedia/what-is-llm-security</a>	T1
<b>LLM Security: Protecting Models, RAG &amp; Data Pipelines - Wiz</b>	<a href="https://www.wiz.io/academy/ai-security/llm-security">https://www.wiz.io/academy/ai-security/llm-security</a>	T1

### DISCLAIMER

This intelligence report is produced by Tech Jacks Solutions Security Command Center (SCC) for informational purposes only. It does not constitute professional security advice, legal counsel, or an incident response engagement. The information herein is derived from publicly available sources and AI-assisted analysis; while every effort is made to ensure accuracy, Tech Jacks Solutions makes no warranties regarding completeness or timeliness. Organizations should conduct their own validation and consult qualified security professionals before taking action based on this report. Tech Jacks Solutions is not liable for any damages resulting from the use of this information.

Generated 2026-07-02 07:12 UTC by TJS Security Command Center